

Supervisory Control of a Pilot-Scale Cooling Loop

4th International Symposium on Resilient Control Systems

Kris Villez
Venkat Venkatasubramanian
Humberto Garcia
Craig Rieger

August 2011

The INL is a
U.S. Department of Energy
National Laboratory
operated by
Battelle Energy Alliance



This is a preprint of a paper intended for publication in a journal or proceedings. Since changes may be made before publication, this preprint should not be cited or reproduced without permission of the author. This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, or any of their employees, makes any warranty, expressed or implied, or assumes any legal liability or responsibility for any third party's use, or the results of such use, of any information, apparatus, product or process disclosed in this report, or represents that its use by such third party would not infringe privately owned rights. The views expressed in this paper are not necessarily those of the United States Government or the sponsoring agency.

Supervisory control of a pilot-scale cooling loop

Kris Villez, Venkat Venkatasubramanian
School of Chemical Engineering
Purdue University
West Lafayette, IN, USA
{kvillez,venkat}@purdue.edu

Humberto Garcia, Craig Rieger
Idaho National Laboratory
PO Box 1625,
Idaho Falls, ID, USA
{humberto.garcia,craig.rieger}@inl.gov

Abstract — We combine a previously developed strategy for Fault Detection and Identification (FDI) with a supervisory controller in closed loop. The combined method is applied to a model of a pilot-scale cooling loop of a nuclear plant, which includes Kalman filters and a model-based predictive controller as part of normal operation. The system has two valves available for flow control meaning that some redundancy is available. The FDI method is based on likelihood ratios for different fault scenarios which in turn are derived from the application of the Kalman filter. A previously introduced extension of the FDI method is used here to enable detection and identification of non-linear faults like stuck valve problems and proper accounting of the time of fault introduction. The supervisory control system is designed so to take different kinds of actions depending on the status of the fault diagnosis task and on the type of identified fault once diagnosis is complete. Some faults, like sensor bias and drift, are parametric in nature and can be adjusted without need for reconfiguration of the regulatory control system. Other faults, like a stuck valve problem, require reconfiguration of the regulatory control system. The whole strategy is demonstrated for several scenarios.

Keywords—*Fault Detection and Identification, Supervisory control, Kalman filter, model-based diagnosis, Model Predictive Control (MPC)*

I. INTRODUCTION

Fault Detection and Identification (FDI) and supervisory control are loosely coupled research fields with many efforts published in the last two decades. Research in FDI has given rise to many techniques based on different assumptions and philosophies. A general introduction to the field can be found in [1,2,3]. In this paper, the applied FDI strategy belongs to a set of techniques based on quantitative first-principle models. Supervisory control is less broad and more intuitive as a field. In essence, supervisory control is usually set up as a flexible control level which is able to adjust low-level control systems on the basis of available information. Such may be changing performance objectives or awareness of changing process conditions, including results obtained by FDI. Supervisory control may include parametric adjustments such as re-tuning of local controllers or sensor or actuator signal corrections. It also includes structural changes such as making changes in the control structure of a plant. The effective coupling of FDI and supervisory control is still a relatively new area with little real-life experiences published today [4-5]. In this work, we therefore present a strategy in which a Kalman filter based FDI method is coupled with a supervisory controller which

switches between several, pre-designed control strategies. The strategy is tested in silico for a pilot-scale cooling loop built at the Idaho National Laboratory (INL), prior to real-life implementation and fits into a more general research line aimed at the construction of design resilient systems [6].

II. MATERIALS AND METHODS

A. Simulated system

The simulated system is conceptually similar to a pilot-scale Machine Condition Monitoring (MCM) test bed constructed at INL. This pilot-scale plant consists of a water cooling loop (found in nuclear plants) and allows the introduction of several faults such as sensor, valve and pump malfunctions. The structure and parameters were designed so to mimic the behavior of this plant based on initial experiments aimed at system identification. For the purpose of this study, a reduced version of the pilot-scale plant is simulated. This reduced loop consists of a closed loop with one pump and two valves, according to the scheme in Figure 1. The pump works at fixed speed throughout all simulations.




Figure 1. Scheme of the cooling loop without control system

To simulate the flow rate (Q) in the system, the steady state (Q_{ss}) flow is computed algebraically from the vane positions of the two valves (v_1, v_2) as follows:

$$Q_{ss}(t) = Q_{max} \cdot \left[\left(\frac{v_1(t) - v_{1,low}}{v_{1,upp} - v_{1,low}} \right) \cdot \left(\frac{v_2(t) - v_{2,low}}{v_{2,upp} - v_{2,low}} \right) \right]^\alpha \quad (1)$$

This equation expresses several aspects observed with the pilot-scale system. First, both valves shut off the flow rate completely at a degree of opening higher than their minimal position (resp. at $v_{1,low}$, $v_{2,low}$). In addition, above a certain degree of opening, they do not affect the flow rate anymore ($v_{1,upp}$, $v_{2,upp}$). In between these positions, the valve positions have a close to linear effect on the steady state flow rate. The deviation from linearity is introduced by the parameter α . A positive value for α (1.2) is used and results in a higher value for the steady state flow rate compared to a truly linear model. Figure 2 shows the steady state flow rate as function of the two valve vane positions. The dynamic response of the flow rate is modeled as a first order linear response to the steady state flow rate with time constant τ . Mathematically we write:

$$\tau \cdot \frac{dQ(t)}{dt} = Q_{ss}(t) - Q(t) \quad (2)$$

A flow rate measurement (Q_m) is taken every second, subject to random error (e_Q) as follows:

$$Q_m(t) = Q(t) + e_Q(t), e_Q \sim N(0, \sigma_Q) \quad (3)$$

The valve subsystems are both simulated with the same model, yet with different parameters. The valves in the pilot-scale MCM test bed are of the butterfly valve type. For reasons of safety, the motor axle and valve vane axle are not the same. Instead, a system of three axles with two couplings makes sure that disturbances on the valve vane do not transfer completely to the motor and vice versa. To this end, the two couplings leave a gap between axles where they interlock. As a result, when the motor position changes direction, the subsequent intermediate axle only follows after the motor axle has crossed the gap. The same holds for the vane axle when the intermediate axle changes direction. The valve motors have one speed only in both directions (s_1 , s_2). The input signal to each valve (u_1 , u_2) is -1, 0 or 1 for closing, steady and opening behaviors of the motor position. The movement of the motor is delayed (delay time: $t_{d,1}$, $t_{d,2}$). In mathematical terms, the first valve is described as follows. Its motor axle position is described by a simple Ordinary Differential Equation (ODE):

$$\frac{dm_1(t)}{dt} = s_1 \cdot u_1(t - t_{d,1}) \quad (4)$$

For the first valve, the intermediate axle and vane axle position (i_1 , v_1) follow from the following algebraic equations:

$$\begin{aligned} i_1(t) &= \min \left[\max \left(m_1 - g_{m,1}, i_1(t - \delta) \right), m_1 + g_{m,1} \right] \\ v_1(t) &= \min \left[\max \left(i_1 - g_{i,1}, v_1(t - \delta) \right), i_1 + g_{i,1} \right] \end{aligned} \quad (5)$$

where $g_{m,1}$ and $g_{m,2}$ represent the gaps resulting from the respective axle couplings and δ representing an infinitesimally small time step. Intermediate axle position measurements ($i_{1,m}$) are taken each second, subject to random error ($e_{i,1}$):

$$i_{1,m}(t) = i_1(t) + e_{i,1}(t), e_{i,1}(t) \sim N(0, \sigma_{i,1}) \quad (6)$$

One needs only to change the respective valve index in the above equations to obtain the equations for the second valve.

The described valves represent a particular challenge with respect to regulatory control of the whole system for several reasons. First, the gaps due to the axle coupling result in strong hysteresis effects of the backlash type (i.e. piece-wise linear) if classic PID-like control strategies are used. Second, a changing deviation exists between the measurement on the intermediate axle and the vane position which one wants to truly know and control. Depending on whether the intermediate axle moves with the motor, this deviation may range from $-g_{m,1}$ to $+g_{m,1}$. Thirdly, the input signals to the valve are discrete and of limited range. A conventional control strategy will suffer from jitter in the control signal as a result with potential negative effects on the valve conditions. These elements had to be taken into account prior to the design of a supervisory control strategy. All parameters and variables of the simulated system are found in Table 1 together with their description and units.

B. Kalman filter for the valve subsystems

A Kalman filter is set up for each of the valve separately and is executed each second. We describe this here for a single valve using k as the discrete time index. The Kalman filter essentially consists of prediction and updating equations. For the valve system, the motor axle position prediction ($x_{m,p}$) follows from the motor axle position state estimate:

$$x_{m,p}(k+1) = x_{m,s}(k) + s \cdot u \quad (7)$$

The predictions of the intermediate axle and vane positions ($x_{i,p}$, $x_{v,p}$) follow from conditional algebraic equations:

$$\begin{aligned} x_{i,p} &= \min \left[\max \left(x_{m,p}(k) - g_m, x_{i,s}(k-1) \right), x_{m,p}(k) + g_m \right] \\ x_{v,p} &= \min \left[\max \left(x_{i,p}(k) - g_i, x_{v,s}(k-1) \right), x_{i,p}(k) + g_i \right] \end{aligned} \quad (8)$$

with $x_{i,s}$ and $x_{v,s}$ the state estimates. The variance of the motor position prediction (P_p) is computed as:

$$P_p = P_s + G \cdot G' \quad (9)$$

with P_s the state estimate variance and G representing the input noise standard deviation (considered 0 here). We use the same equations for the intermediate axle and vane position variance. This is a naive approach as (1) the intermediate axle and vane positions are correlated when they move along with their coupled axle and uncorrelated when they do not and (2) we assume that the couplings between axles are not a source of noise. However, a complete treatment of the variance-covariance predictions could be based on a hybrid model yet this would lead to a high computational load which is to be avoided for on-line applications. In addition, this naive approach is shown to work well in this work.

Given an intermediate axle position measurement, $y_i(k)$, the state update for the intermediate axle are written as follows:

$$\begin{aligned} r_i(k) &= y_i(k) - x_{i,p}(k) \\ x_{i,s}(k+1) &= x_{i,p}(k) + K \cdot r_i(k) \\ P_s &= P_p - K \cdot \sigma_i^2 \cdot K' \\ K &= P_s \cdot 1 / \sigma_i^2 \end{aligned} \quad (11)$$

with r_i the prediction error, K the Kalman gain and σ_i the measurement variance. Once more, the same equations are used to update the motor and vane position estimates:

$$\begin{aligned} x_{m,s}(k+1) &= x_{m,p}(k) + K \cdot r_i(k) \\ x_{v,s}(k+1) &= x_{v,p}(k) + K \cdot r_i(k) \end{aligned} \quad (12)$$

This is again considered a naive approach as the obtained measurement only reflects on the position of all axles if all axles are moving along with each other. The above equations are executed each time a new measurement is available.

C. Regulatory control system

A control system is setup as follows. A master controller defines valve movement based on measurement of the flow rate and a model of the flow rate as function of the valve vane positions. The master controller makes use of the single valve indicated by the supervisory controller. This will be valve 1 in normal situations as it is the fastest valve and has the most accurate measurement. The controller computes the expected flow rate (Q_p) for the current valve positions first, according to the equation 1 with α set to 1, thus assuming a linear, static approximation to the real system. A prediction error (r_Q) is the difference between expected flow rate and measurement:

$$r_Q(k) = Q_p(k) - Q_m(k) \quad (13)$$

The prediction error is summed into an integral term at each time step (i_Q), similar to the integrative part of a conventional discrete-time PID controller. With τ_i the integral time constant:

$$i_Q(k) = i_Q(k-1) + r_Q(k) / \tau_i \quad (14)$$

Next, one computes the desired valve position by means of inversion of equation 1 with α set to 1 and the other valve position set to its current estimate ($x_{v,s,2}$) and for a flow rate set to $Q_{sp}(k) + i_Q(k)$ with $Q_{sp}(k)$ the actual setpoint. By means of the correction term $i_Q(k)$, the effect of system-model mismatch is reduced. We write the inversion generally as

$$v_{sp,1} = f^{-1}(x_{v,s,2}, Q_{sp} + i_Q(k)) \quad (15)$$

with f^{-1} meaning the inversion of equation 1. The desired valve setpoint is now compared to the predicted valve position, taking delays of the control signal to the first valve into account. To this end, the prediction equations of the Kalman

filter are iterated to match the time delay. Following this, one compares the valve setpoint with the predicted vane position. If the difference is large enough, the control signal is set to +1 or -1 according to the desired direction of movement. Otherwise, the valve control signal is set to 0. The difference between setpoint and predicted position is considered high enough if the setpoint lies outside of the 3- σ confidence limits (approx. 99% confidence limit) of the valve vane position prediction. By this logic, one can avoid jitter in the control signal and account for changing precision of the valve vane position prediction.

D. Simulated scenarios

A total of five scenarios is simulated. The first scenario is faultless. The initial valves are at 100% opening and at the start of the simulation, the setpoint is set to 100 gpm. The total simulated time is 200 seconds. All following scenarios are the same except for added faults. All faults are introduced at 50 seconds in the simulation. In the second scenario the first valve gets stuck. In the third scenario, a bias of 10% is introduced in the first valve measurement. A fourth consists of a drift in the first valve measurement with a speed of 0.5% per second. The fifth scenario consists of both valves being stuck. With this set of simulations, a perfect match exists between simulated faults and faults available in the fault diagnosis module. While this is widely recognized as a potential drawback we do not investigate this particular aspect of the method in this work.

E. Fault detection and diagnosis

The fault detection and diagnosis method as described in [4] is applied. This strategy was originally set up for linear models and linear, additive faults and is based on log-likelihood ratios following fault parameter estimation. In [7] this strategy was successfully expanded for other types of faults including stuck, stiction and drift and was improved by considering the time of fault introduction as a fault parameter. In the MCM case, excessive friction will lead to a stuck valve fault rather than valve stiction. As such, this type of fault is not be included here. We do include the stuck and drift problems however and proper identification of fault time introduction. We stress the essential elements of the method here.

The first step in the FDI method is fault detection. Fault detection follows from comparing the sum of squared prediction residuals (SSPR) by the Kalman filter to the 99% confidence level limit of the corresponding χ^2 distribution, assuming normal operation of the system. A detection occurs when the SSPR crosses the 99% confidence limit. In our case, a fault detection is confirmed when 7 tests are positive within the last 10 consecutive prediction steps.

Following a confirmed fault detection step, the FDI module collects data of 20 more data samples. Following this, diagnosis starts. To this end, the last 60 data samples are used (40 before fault confirmation, 20 after). Each time instant within this window is considered as a candidate value for the time of fault introduction. For each of these (60) time instants, the maximum likelihood for each fault type is evaluated. For the stuck valve problem, there are no other parameters than the time of introduction. In contrast, the valve measurement bias and drift contain one parameter each which is identified for

each considered time of fault introduction. This is a linear operation, as pointed out in [4,5,7] and can thus be obtained very fast. Following these identifications, one has obtained the maximum likelihood for each fault type conditional to the considered fault introduction time. One selects the scenario with maximum likelihood along fault introduction time variable and fault type, thereby completing the FDI step.

F. Supervisory controller

A supervisory controller is set up to accommodate results from fault detection and from fault diagnosis. The supervisory controller keeps the control strategy with the first valve as long as no fault is confirmed. When a fault is confirmed, a safety mode is activated in which both valves are opened. This is considered a safe operation in a cooling context as it results in maximum cooling capacity. This safety mode is held as long as no conclusive fault diagnosis is given. This also results in movement of the valve -except when it is stuck- and assists in obtaining informative data about the nature of the fault [7].

Following fault diagnosis, several actions can be taken. For bias and drift faults, a so-called parametric adjustment is executed. That is to say that the respective measurement is corrected for the identified fault. E.g. the identified bias is subtracted from the received measurement. Following this, the control strategy which used the first valve to control the flow is reinstated, bringing the system back to normal operation. We note that the method also allows automatic adjustment if the bias or drift are removed again. Indeed, the removal of bias or accumulated drift will appear as a bias which can be identified again and adjusted for as demonstrated in [4]. For a stuck valve, such parametric adjustment is not sufficient since the first valve is now useless for control of the flow rate. For this reason, another control strategy is activated in which the second, less performing valve (slower, less accurate), is used to control the flow. To compute the second valve position setpoint, the first valve position is assumed to be the identified stuck position. This action by the supervisory controller is structural in nature and allows to mitigate a complete failure of the first valve. In such case, it is possible to obtain faults in the second valve as well, both parametric faults or complete failure. While the parametric faults can still be mitigated, a complete failure of both valves cannot be mitigated with the proposed supervisory controller. In such an event, the supervisory controller maintains the safety mode and alarms the operator or systems higher in the control hierarchy. We note that the developed method does not include an automated strategy to go back to the original control strategy if the first valve is functional again, e.g. following maintenance.

G. Performance evaluation and comparison

The proposed supervisory control system is evaluated by means of the Integral Time Absolute Error (ITAE) computed over the simulated time in each of the simulated scenarios. This ITAE is also evaluated for two other strategies which are reduced forms of the proposed one. A first alternative consists of the complete supervisory controller except for structural adjustments, i.e. as if the second valve does not exist. A second alternative system consists of the FDI method and includes the safety mode controller. Irrespective of fault

identification, the first valve is kept in use, i.e. no parametric or structural adjustment is made to the regulatory control system.

III. RESULTS

A. Scenario 1: Faultless system

Figure 2 shows the results for a normal, faultless simulation of the system. The valve position measurements can hardly be discerned from the true values and as is the case for the confidence bounds on valve position estimates by the Kalman filter after just a few seconds. Furthermore, a true flow of 102.5 gpm is reached at about 70 seconds in the simulation by manipulation of valve 1 and remains thereafter. This offset is due to the jitter-avoiding control strategy as explained in section 2.C. Valve 2 remains completely open.




Figure 2. Scenario 1 - Faultless operation. Cross-hairs indicate measurements. Full lines indicate true value and dashed lines indicate 5-sigma confidence bounds of the Kalman state estimate. Top: valve 1 intermediate axle; middle: valve 2 intermediate axle; bottom: flow rate.




Figure 3. Scenario 2 – Valve 1 stuck. Cross-hairs indicate measurements. Full lines indicate true value and dashed lines indicate 5-sigma confidence bounds of the Kalman state estimate. Vertical lines indicate the times of fault confirmation (full) and fault diagnosis (dashed). Top: valve 1 intermediate axle; middle: valve 2 intermediate axle; bottom: flow rate.

B. Scenario 2: Valve 1 stuck

Figure 3 shows results for the second scenario, in which the first valve gets stuck at 50 seconds in the simulation. At 61 seconds, this fault is confirmed. One observes the deviation between measurements and Kalman confidence bounds clearly. The supervisory controller switches to safety mode instructing

both valves to open. Since the first valve fails to do so, the fault is correctly identified easily at 81 seconds. Following this correct identification, the supervisory controller switches to the control strategy using the second valve. Measurements of valve 1 are discarded beyond this point. A flow of 112 gpm is now reached at about 170 seconds in the simulation. The offset is larger as the jitter-avoiding controller deals with a less precise position estimate, due to larger noise in the valve measurement.

C. Scenario 3: Valve 1 bias

The results for the valve measurement bias are shown in Figure 4. Here, the fault is detected at time 57 seconds and correctly identified at 77 seconds. In between these two steps, valve 1 opens following the switch to a safety control strategy. After fault identification, the measurement fault is corrected for and the normal control strategy is resumed. A flow of 102.45 gpm is reached, similar to the faultless scenario at 130 seconds in the simulation. Obviously, the safety procedure which opens valve 1 delays the time by which the setpoint is reached.




Figure 4. Scenario 3 – Valve 1 bias. Cross-hairs indicate measurements. Full lines indicate true value and dashed lines indicate 5-sigma confidence bounds of the Kalman state estimate. Vertical lines indicate the times of fault confirmation (full) and fault diagnosis (dashed). Top: valve 1 intermediate axle; middle: valve 2 intermediate axle; bottom: flow rate.

D. Scenario 4: Valve 1 drift

The valve drift scenario results are shown in Figure 5. Here, the fault is confirmed at 66 seconds, followed by correct identification at 86 seconds. As with the bias fault, the valve opens during the period in between due to the activated safety control strategy. Following fault identification, a parametric correction is executed and the original control strategy is resumed. One can observe that the correction is not perfect as the measurement remains to drift, yet at a slower pace. This is because the drift slope is harder to estimate than a bias. The remaining drift is however not large enough to trigger a detection. A flow of 102.5 gpm is reached at about 110 seconds and then remains. Thus, despite the imprecision of the drift parameter estimate, the setpoint can be reached.

E. Scenario 5: Valve 1 and valve 2 stuck

Figure 6 shows the results for the scenario in which both valves get stuck. The stuck valve problem for valve 1 is recognized first at 62 seconds, followed by correct identification at 82 seconds. Then, the supervisory controller

switches to the second control strategy, thereby trying to use valve 2 for flow control. This also fails as it is stuck as well. At 94 seconds it is confirmed and at correctly identified at 114 seconds. Now, the supervisory controller cannot switch to any viable control strategy. It therefore remains in safety mode and alarms the higher level in the control structure (not shown).




Figure 5. Scenario 4 - Valve 1 drift. Cross-hairs indicate measurements. Full lines indicate true value and dashed lines indicate 5-sigma confidence bounds of the Kalman state estimate. Vertical lines indicate the times of fault confirmation (full) and fault diagnosis (dashed). Top: valve 1 intermediate axle; middle: valve 2 intermediate axle; bottom: flow rate.




Figure 6. Scenario 5 - Valve 1 and 2 stuck. Cross-hairs indicate measurements. Full lines indicate true value and dashed lines indicate 5-sigma confidence bounds of the Kalman state estimate. Vertical lines indicate the times of fault confirmation (full) and fault diagnosis (dashed). Top: valve 1 intermediate axle; middle: valve 2 intermediate axle; bottom: flow rate.

F. Comparison with reduced supervisor alternatives

We compare the proposed supervisory control strategy with two reduced alternatives for the 5 simulated scenarios. We refer to the complete strategy as strategy I. The strategy without structural adjustments is called strategy II and the strategy without supervisory control as strategy III. Table 2 shows the resulting average Integral Time Absolute Error (ITAE) computed on the basis of the measured flow rate. In scenario 1 (no fault), there is no effect of the supervisory control choice on the ITAE, as expected. For scenario 2, a positive effect is seen when choosing the complete supervisory control system (strategy I) over strategy II and III. This is again expected as a stuck valve can only be mitigated with a structural adjustment. For scenarios 3 and 4, it can be seen that either strategy I or II deliver the same, optimal, performance, as they both permit the

desired parametric adjustments. Finally, for scenario 5, none of the alternative supervisory control strategies is able to mitigate both valves being stuck, resulting in the highest ITAE values across the table.

TABLE I. INTEGRAL TIME ABSOLUTE ERROR (ITAE) FOR DIFFERENT SCENARIOS AND SUPERVISORY CONTROL STRATEGIES.

Strategy	Scenario				
	1	2	3	4	5
I	45.55	81.41	61.07	55.23	98.78
II	45.55	98.78	61.07	55.23	98.78
III	45.55	98.78	61.51	61.82	98.78

CONCLUSIONS

In this paper, a supervisory control strategy is evaluated in silico prior to on-line application with a pilot-scale plant mimicking a secondary cooling loop of a nuclear power plant. The supervisory control system aims at the proper choice between a set of control strategies based on the continuous assessment of the condition of each of the available valves. A Kalman-based strategy is used for timely fault detection and diagnosis of the valve subsystems. Following this identification, the supervisory controller permits both parametric adjustments as well as structural changes in the regulatory control structure. The complete strategy enables to respond adequately to a variety of disturbances such as noise, faults and failures thereby contributing to improved resilience of the studied system. As such, conventional control, parametric adjustments and At the same time, the limitations of the supervisory controller are shown in a scenario in which both valves get stuck and no further structural changes with the given actuating and sensing equipment is possible.

ACKNOWLEDGMENT

The authors wish to thank the ICIS Distinctive Signature at Idaho National Laboratory (INL) for the support of this work.

REFERENCES

- [1] Venkatasubramanian, V., Rengaswamy, R., and Kavuri, S. N., "A review of process fault detection and diagnosis - Part I: Quantitative model-based methods". Comput. Chem. Eng., 2003, 27, 293-311.
- [2] Venkatasubramanian, V., Rengaswamy, R., Yin, K., and Kavuri, S. N., "A review of process fault detection and diagnosis - Part II: Qualitative models and search strategies." Comput. Chem. Eng., 2003, 27, 313-326.
- [3] Venkatasubramanian, V., Rengaswamy, R., Kavuri, S. N., and Yin, K., "A review of process fault detection and diagnosis - Part III: Process history based methods." Comput. Chem. Eng., 2003, 27, 327-346.
- [4] Prakash, J., Narasimhan, S., and Patwardhan, S. C., "Integrating model based fault diagnosis with model predictive control." Ind. Eng. Chem. Res., 2005, 44, 4344-4360.
- [5] Prakash, J., Patwardhan, S. C., and Narasimhan, S., "A supervisory approach to fault-tolerant control of linear multivariable systems." Ind. Eng. Chem. Res., 2002, 41, 2270-2281.
- [6] Rieger, C. G., Gertman, D. I., and McQueen, M. A., "Resilient control systems: Next generation design research." In 2nd IEEE Conf. on Human System Interaction, Catania, Italy, May 2009.
- [7] Villez, K., Srinivasan, B., Rengaswamy, R. and Narasimhan, S., "Kalman-based strategies for Fault Detection and Diagnosis: Extensions

and critical evaluation for a simple benchmark system." Comp. Chem. Eng., 35, 806-816.

TABLE II. SIMULATED MODEL PARAMETERS AND VARIABLES

Symbol	Description	Value	Unit
<i>Model parameters</i>			
α	Flow rate exponent	1.2	[-]
τ	Flow rate time constant	2	[s]
$\sigma_{i,1}, \sigma_{i,2}, \sigma_Q$	Noise standard deviations	1, 2, 5	[gpm, %, %]
$g_{m,1}, g_{m,2}$	Gaps at motor and intermediate axle	1, 1	[%]
$g_{i,1}, g_{i,2}$	Gaps at intermediate and vane axle	2, 2	[%]
s_1, s_2	Valve speeds	1, 9.8	[%/s]
$v_{1,low}, v_{2,low}$	Valve lower effect limits	20, 10	[%]
$v_{1,upp}, v_{2,upp}$	Valve upper effect limits	70, 20	[%]
$t_{d,1}, t_{d,2}$	Valve delays	2, 2	[s]
Q_{max}	Maximum flow rate	240	[gpm]
<i>Model variables</i>			
$e_{i,1}, e_{i,2}, e_Q$	Measurement error in valve position and flow measurements		[%]
i_1, i_2	Valve intermediate axle positions		[%]
m_1, m_2	Valve motor axle positions		[%]
u_1, u_2	Valve input signals		[-]
v_1, v_2	Valve vane positions		[%]
Q, Q_{ss}, Q_m	Flow rate: actual, steady state and measurement value		[gpm]
<i>Kalman filters and regulatory control system</i>			
σ_i	Intermediate axle noise standard deviation		[%]
τ_i	Integrator time constant for flow control		[s]
i_Q	Flow rate integrated prediction error		[gpm]
r_Q	Flow rate prediction error		[gpm]
r_i	Intermediate axle prediction error		[%]
s	Valve speed		[%/s]
u	Valve input signal		[-]
$v_{SP,1}, v_{SP,2}$	Vane position setpoint		[gpm]
$x_{i,p}, x_{i,s}$	Intermediate axle state prediction / estimate		[%]
$x_{v,p}, x_{v,s}$	Vane axle state prediction / estimate		[%]
$x_{m,p}, x_{m,s}$	Motor axle state prediction / estimate		[%]
y_i	Intermediate axle measurement		[%]
g_m, g_i	Gaps at axle couplings, assumed known		[%]
G	Input noise variance		[%]
K	Kalman gain		[-]
P_p, P_s	Variance of state predictions / estimates		[% ²]
Q_p, Q_{SP}	Flow rate prediction / setpoint		[gpm]